IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	Eric R. Schott		
Application No.:	10/761,884	Group:	2188
Filed:	January 20, 2004	Examiner:	Duc T. Doan
Confirmation No:	4233		
For:	Storage Systems Having Differentiated Storage Pools		

CERTIFICATE OF MAILING OR TRANSMISSION

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as First Class Mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, or is being facsimile transmitted to the United States Patent and Trademark Office on:

6/29/09 Nadine Z. Kush
Date Signature

Nadine Z. Kush
Typed or printed name of person signing certificate

AMENDMENT

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Amendment is being filed in response to the Office Action mailed from the U.S. Patent and Trademark Office on March 27, 2009 in the above-identified application. Reconsideration and further examination are requested.

Please amend the application as follows:

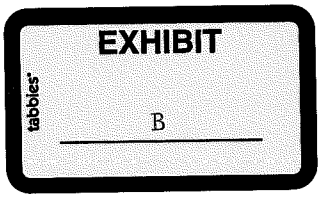

Amendments to the Claims

Please cancel Claims 3 and 13. Please amend Claim(s) 1 and 19. The Claim Listing below will replace all prior versions of the claims in the application:

Claim Listing

1. (Currently Amended) A system for providing differentiated classes of storage, comprising:

a storage device having a plurality of storage locations and a logical block name space for organizing logical block names of the storage locations,

a performance process configured to determine a level of performance for the plurality of storage locations and partition the plurality of storage locations into a plurality of regions as determined by their different levels of performance,

a mapping process configured to map the partitioned regions of the storage locations and aggregate the logical block names of the storage locations of the storage device in the partitioned regions having an identical level of performance to a selected section of the logical block name space,

a RAID controller, for assigning a first RAID level configuration to a first set of aggregated logical block names of the storage device, and assigning a second RAID level configuration to a second set of aggregated logical block names of the storage device, the first and second RAID level configurations being different from each other; and

the system thereby providing differentiated classes of storage having two or more differentiated RAID level configurations of the same storage device to one or more clients accessing the system.

2. (Previously Presented) A system according to claim 1, wherein:

the performance process separates the plurality of storage locations into a plurality of categories as determined by their different levels of performance.

3. (Canceled)

4.-7. (Canceled).

8. (Previously Presented) A system according to claim 1, further comprising:

a process configured to employ the storage system to provide a file system service.

9. (Previously Presented) A system according to claim 1, further comprising:

a process configured to provide a storage volume service.

10. (Previously Presented) A system according to claim 9, wherein the mapping process creates multiple storage volumes at a selected level of performance on the storage device.

11. (Previously Presented) A performance process for providing differentiated classes of storage based on determined levels of performance of a plurality of storage locations of a storage device, the process comprising the steps of:

providing a storage device having a plurality of storage locations and a logical block name space for organizing logical block names of the storage locations,

determining a level of performance of the plurality of storage locations,

partitioning the plurality of storage locations into a plurality of regions as determined by their different levels of performance,

mapping partitioned regions of the storage locations,

aggregating the logical block names of the storage locations in the partitioned regions having an identical level of performance to a selected section of the logical block name space,

assigning a first RAID level configuration to a first set of aggregated logical block names,

assigning a second RAID level configuration to a second set of aggregated logical block names, the first RAID level configuration and second RAID level configuration being different from one another, and

the storage device thereby providing differentiated classes of RAID level storage to one or more clients.

12. (Previously Presented) A process according to claim 11, further including the step of separating the plurality of storage locations into a plurality of categories as determined by their different levels of performance.

13. (Canceled).

14.-17. (Canceled).

18. (Original) A process according to claim 11, wherein mapping creates multiple storage volumes at a selected level of performance.

19. (Currently Amended) A system for providing differentiated classes of storage, comprising

a storage device having a plurality of storage locations, a logical block name space for organizing logical block names of the storage locations, and performance parameters of the storage locations that vary across the storage device;

a partitioning process configured to partition the storage locations into regions and aggregate the logical block names of the storage locations in the partitioned regions having an identical level of performance to a selected section of the logical block name space, thereby providing two or more differentiated classes of storage <u>on the storage device</u> to one or more clients accessing the system;

a RAID controller, for assigning different RAID level techniques to respective ones of the two or more differentiated classes of storage; and

a performance measurement system that scans storage locations of the storage device and determines the level of performance of the storage locations.

20. (Previously Presented) A system according to claim 19, wherein the partitioning process selects a fixed set of partitions as a function of a selected configuration of system components.

21. (Canceled).

22. (Previously Presented) The system of claim 1, wherein a level of performance includes a data access time, or a reliability of a storage location, or a combination thereof.

23. (Previously Presented) The system of claim 1, wherein the storage device is a single storage disk.

24. (Previously Presented) The system of claim 1, wherein the mapping process performs mapping and aggregating when the storage system is designed.

25. (Previously Presented) The system of claim 1, wherein the mapping process performs mapping and aggregating during operation of the storage device.

26. (Previously Presented) The system of claim 1, further comprising a performance measurement system configured to scan storage locations of the storage device and determine the level of performance for the storage locations.

27. (Previously Presented) The system of claim 26, wherein the performance measurement system performs experimental read and write operations and determines the level of performance from experimental data collected in the read and write operations.

28. (Previously Presented) The process of claim 11, wherein a level of performance includes a data access time, or a reliability of a storage location, or a combination thereof.

29. (Previously Presented) The process of claim 11, wherein partitioning comprises performing experimental read and write operations and determining the level of performance from experimental data collected in the read and write operations.

30. (Previously Presented) The process of claim 11, wherein mapping and aggregating are performed when a storage system that implements the process is designed.

31. (Previously Presented) The process of claim 11, wherein mapping and aggregating are performed during operation of a storage system that implements the process.

REMARKS

With entry of the foregoing amendment claims 1-2, 8-12, 18-20 and 22-31 remain in the application.

Claims 1and 19 are now being amended. Support for these amendments can be found in the specification and drawings as originally filed. For example, Figs. 3 and 4 of the specification and the accompanying text at pages 8 through 11 explain how differentiated storage pools are created by assigning different RAID configuration levels to the storage device locations on a single disk device that are grouped according to performance variations. Fig. 4 depicts a specific example of an extent based (or page-based) allocation operation where the differentiated storage pools are of the type depicted in Fig. 3. In a particular example, the allocation operation provides a first pool A of the available storage to support a first RAID level 10. A second pool B of storage subspace is employed to support a RAID 5 service level, and a subspace C is configured to support a RAID 50 performance level.

Power of Attorney

Acknowledgement of the Power of Attorney filed March 24, 2008 and correction of the correspondence address in the Patent Office's records is respectfully requested.

Claim Objections

The Applicant appreciates the acknowledgement that the objection to Claim 11 has been withdrawn.

Specification Objections

The last Office Action, at the top of page 11, indicates that Applicant's remarks in the response filed February 10, 2009 overcame the claims objection. However, there still is no clear statement on the record that the previous response overcame the objection to the specification, specifically the previous objection to the priority claim to U.S. Provisional Application Serial

No. 60/441,810 filed January 21, 2008. Clarification that all objections to the priority claim and any other aspect of the specification has been withdrawn is again requested.

The Claims as Amended are not Rendered Obvious by Dimitri in view of Jacobsen and the PC Guide

Claims 1-2, 8-12, 18, and 22-31 stand rejected under 35 U.S.C. 103 as being obvious. With entry of the foregoing amendment, claims 1, 11, and 19 now all more particularly recite that the system or process has several differentiated performance Logical Block Name (LBN) subspaces which respectively support a different one of several RAID levels on a single storage device. This feature not found in the prior art.

As one example of the invention described in the application, a single RAID system provides different levels of RAID storage functionality for different classes, or "pools" of storage. A measurement system scans individual devices (such as a depicted storage device 10 of Fig. 1 of the application), and measures certain characteristics to determine how the device should be subdivided. Using the results of this measurement process, a determination is made that a collection of logical block names (LBNs) can be aggregated. The aggregated blocks having a common performance characteristics then provide a subspace within the LBN space for that single device. See page 7 line 9 through page 8 line 9 of the utility application as originally filed.

Applicants' system 20 depicted in Fig. 2 can employ multiple storage devices to support two different classes of storage (in a general sense), but specifically supports two different RAID levels on each storage device. In this particular example, the same storage device is supporting both RAID level 10 with pool A22 and RAID level 50 with pool B24.

Fig. 4 and the accompanying text at page 10 line 11 through page 11 line 5 illustrate this advantage. For example, a given device may have two areas that are (1) a high activity area assigned to pool A and other areas that are (2) low activity areas (highlighted in white) that are assigned to another pool C. Pool A areas may have high activity thus maybe assigned to a RAID 10 array. The RAID 10 array provides higher performance for higher activity but lower

redundancy. At the same time, low activity areas assigned to pool C can be used to support a RAID 50 configuration. These slower performing areas can provide higher redundancy.

Prior to the Applicants' invention, systems such as that shown in Dimitri and Jacobsen and PC Guide would have required different physical devices to support different RAID level types. Applicants do note that Dimitri does provide a system that enables differentiated "classes" of storage and organizes these according to performance levels. (These are called "zones", since Dimitri's disk is a zone constant angular velocity ("VCAV") formatted disk.) A VCAV disk media inherently spins at a constant rotation per minute such that inner zones inherently have a greater data rate per revolution than outer zones. While it is not clear how logical block addresses (LBAs) map to "zones" in Dimitri, it is clear that his system will select a zone according to a set of performance criteria and treat them similarly. For example, at Fig. 4 and column 8 lines 31 through 43, Dimitri does teach that a zone would be selected based on a combination of file size and utilization history of the file. After selecting a zone, the zone control unit 50 then writes all data stripes to the same zone on the different disks.

The Dimitri architecture thus avoids a situation where overall performance of a RAID array is limited to the performance of the inner most zone of the disks.

However, Dimitri only mentions RAID generally and does not recognize the possibility of different RAID levels at all. In fact, all that Dimitri ever suggests using is a type of disk stripping which equates to RAID level 0. (Dimitri at column 8 lines 22 through 25.)

Especially with entry of the foregoing amendment, Dimitri does not provide or even suggest a way to support different RAID levels. For this reason alone, claim 1 should be allowed.

In other aspects of the present invention, a performance process measures the performance of storing locations by making experimental read and write operations across the logical block main space, and uses these measurements to determine whether various locations can be aggregated to regions. (See for example, Applicants' specification at page 7, lines 4-11). A mapping process thus aggregates the logical block names for location having an identical level of performance (i.e., the partition locations). This results in the creation of different storage pools. The Applicants' claimed system clearly assigns different RAID levels to these different regions of a device based upon the determined levels of performance of the locations within the

regions (see at least the Applicants' specification at page 2, lines 17 through 19, and page 4 lines 9 through 11, and Fig. 4). Clients thus accessing such a system can utilize the storage pool and select from a number of different classes of RAID level service from the very same storage array (see the specification at page 11, lines 5 through 8). For example, one client may utilize a RAID 10 level service, while another client may utilize a RAID 5 level service, both however storing their data on the same set of physical disks (see the specification at page 9, line 22 through page 10 line 2). This is not taught or suggested by Dimitri.

The Examiner is also incorrect in his assessment that Jacobsen discloses a RAID controller for assigning a first RAID level configuration to a first set of aggregated logical block names.

Jacobsen does show a disk array controller that can control a first set of disks (26, 28) to provide mirror level redundancy and a second set of disks (24, 30 and 32) that can be used to provide parity level redundancy. But that is it. Jacobsen mentions nothing at all whatsoever concerning logical block names (LBN). The Examiner's conclusions about Jacobsen are erroneous for that reason alone.

Jacobsen also only discusses the configuration of groups of disks to provide a certain level of RAID performance. For example, multiple disks 26, 28 are providing mirror type functionality and multiple disks 24, 30 and 32 are providing parity level performance. There is no suggesting, teaching, or any inference in Jacobsen that any particular one of the disks can provide two different RAID levels of performance.

For these reasons, the Examiner's rejection of the Applicant's claims 1, 11 and 19 is deficient as failing to set forth a prima facie case of obviousness.

As admitted by the Examiner, while the PC Guide does further disclose assigning two RAID levels to a set of disks, it does not teach, suggest or infer that those two RAID levels can be provided by a single device.

All other claims depend from claim 1, 11, or 19.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By_____

David J. Thibodeau, Jr.
Registration No. 31,671
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Date: 6/29/09